UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/826,271	04/04/2001	Johannes Markus Hoerler	112025-0466	3392

24267 7590 09/08/2004

CESARI AND MCKENNA, LLP
88 BLACK FALCON AVENUE
BOSTON, MA 02210

EXAMINER	
KING, JUSTIN	

ART UNIT	PAPER NUMBER
2111	

DATE MAILED: 09/08/2004

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	09/826,271	HOERLER ET AL.	
	Examiner	Art Unit	
	Justin I. King	2111	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).
- Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *18 August 2004* .

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-37* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-37* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

11)☐ The proposed drawing correction filed on _____ is: a)☐ approved b)☐ disapproved by the Examiner.

If approved, corrected drawings are required in reply to this Office action.

12)☐ The oath or declaration is objected to by the Examiner.

Priority under 35 U.S.C. §§ 119 and 120

13)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____ .

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

14)☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. § 119(e) (to a provisional application).

a) ☐ The translation of the foreign language provisional application has been received.

15)☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. §§ 120 and/or 121.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO-1449) Paper No(s) _____ .
4) ☐ Interview Summary (PTO-413) Paper No(s). _____ .
5) ☐ Notice of Informal Patent Application (PTO-152)
6) ☐ Other: .

DETAILED ACTION

Claim Rejections - 35 USC § 102

1. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (a) the invention was known or used by others in this country, or patented or described in a printed publication in this or a foreign country, before the invention thereof by the applicant for a patent.

2. Claims 1, 7, 11, 19-21, and 35-37 are rejected under 35 U.S.C. 102(a) as being anticipated by Greim et al. (U.S. Patent No. 6,163,829).

Referring to claim 1: Greim discloses an external device (figure 4, structure 18 and the attached device from path 12) coupled to a high latency path (figure 4, structure 12), the external device generating a pulsed interrupt signal for each type of interrupt supported by the processor (figure 4, structure 20); an interrupt multiplexing device (figure 4, structure 82) accessible by the processor over a fast bus (figure 4, combined structures of 80 and 22), the interrupt multiplexing device adapted to issue the interrupt to the processor in response to each pulsed interrupt signal generated by the external device (column 2, first paragraph); a low latency path (figure 4, dedicated bus between structures 18 and 82) coupling the external device to the interrupt multiplexing device and adapted to transport each pulsed interrupt signal generated by the external device to the interrupt multiplexing device; and a status bit (figure 6, structure 142) stored within the interrupt multiplexing device, the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor

efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device over the fast bus. Hence, claim is anticipated by the Greim.

Referring to claim 7: Greim's disclosure is stated above; furthermore, Greim discloses the DMA as the interrupt source (column 24, line 45).

Referring to claim 11: Greim discloses an external device (figure 4, structure 18 and the attached device from path 12) coupled to a high latency path (figure 4, structure 12), the external device generating a pulsed interrupt signal for each type of interrupt supported by the processor (figure 4, structure 20); an interrupt multiplexing device (figure 4, structure 82) accessible by the processor over a fast bus (figure 4, combined structures of 80 and 22), the interrupt multiplexing device adapted to issue the interrupt to the processor in response to each pulsed interrupt signal generated by the external device (column 2, first paragraph); a low latency path (figure 4, dedicated bus between structures 18 and 82) coupling the external device to the interrupt multiplexing device and adapted to transport each pulsed interrupt signal generated by the external device to the interrupt multiplexing device; and a status bit (figure 6, structure 142) stored within the interrupt multiplexing device, the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device over the fast bus. Hence, claim is anticipated by the Greim.

Referring to claim 19: Greim discloses an external device (figure 4, structure 18 and the attached device from path 12) coupled to a high latency path (figure 4, structure 12), the external device generating a pulsed interrupt signal for each type of interrupt supported by the processor (figure 4, structure 20); an interrupt multiplexing device (figure 4, structure 82) accessible by the

processor over a fast bus (figure 4, combined structures of 80 and 22), the interrupt multiplexing device adapted to issue the interrupt to the processor in response to each pulsed interrupt signal generated by the external device (column 2, first paragraph); a low latency path (figure 4, dedicated bus between structures 18 and 82) coupling the external device to the interrupt multiplexing device and adapted to transport each pulsed interrupt signal generated by the external device to the interrupt multiplexing device; and a status bit (figure 6, structure 142) stored within the interrupt multiplexing device, the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device over the fast bus. Hence, claim is anticipated by the Greim.

Referring to claim 20: Greim discloses an external device (figure 4, structure 18 and the attached device from path 12) coupled to a high latency path (figure 4, structure 12), the external device generating a pulsed interrupt signal for each type of interrupt supported by the processor (figure 4, structure 20); an interrupt multiplexing device (figure 4, structure 82) accessible by the processor over a fast bus (figure 4, combined structures of 80 and 22), the interrupt multiplexing device adapted to issue the interrupt to the processor in response to each pulsed interrupt signal generated by the external device (column 2, first paragraph); a low latency path (figure 4, dedicated bus between structures 18 and 82) coupling the external device to the interrupt multiplexing device and adapted to transport each pulsed interrupt signal generated by the external device to the interrupt multiplexing device; and a status bit (figure 6, structure 142) stored within the interrupt multiplexing device, the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor

efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device over the fast bus. Hence, claim is anticipated by the Greim.

Referring to claim 21: Greim discloses an external device (figure 4, structure 18 and the attached device from path 12) coupled to a high latency path (figure 4, structure 12), the external device generating a pulsed interrupt signal for each type of interrupt supported by the processor (figure 4, structure 20); an interrupt multiplexing device (figure 4, structure 82) accessible by the processor over a fast bus (figure 4, combined structures of 80 and 22), the interrupt multiplexing device adapted to issue the interrupt to the processor in response to each pulsed interrupt signal generated by the external device (column 2, first paragraph); a low latency path (figure 4, dedicated bus between structures 18 and 82) coupling the external device to the interrupt multiplexing device and adapted to transport each pulsed interrupt signal generated by the external device to the interrupt multiplexing device; and a status bit (figure 6, structure 142) stored within the interrupt multiplexing device, the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device over the fast bus. Hence, claim is anticipated by the Greim.

Referring to claim 35: Greim discloses an external device (figure 4, structure 18 and the attached device from path 12) coupled to a high latency path (figure 4, structure 12), the external device generating a pulsed interrupt signal for each type of interrupt supported by the processor (figure 4, structure 20); an interrupt multiplexing device (figure 4, structure 82) accessible by the processor over a fast bus (figure 4, combined structures of 80 and 22), the interrupt multiplexing device adapted to issue the interrupt to the processor in response to each pulsed interrupt signal

generated by the external device (column 2, first paragraph); a low latency path (figure 4, dedicated bus between structures 18 and 82) coupling the external device to the interrupt multiplexing device and adapted to transport each pulsed interrupt signal generated by the external device to the interrupt multiplexing device; and a status bit (figure 6, structure 142) stored within the interrupt multiplexing device, the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device over the fast bus. Hence, claim is anticipated by the Greim.

Referring to claim 36: Greim discloses an external device (figure 4, structure 18 and the attached device from path 12) coupled to a high latency path (figure 4, structure 12), the external device generating a pulsed interrupt signal for each type of interrupt supported by the processor (figure 4, structure 20); an interrupt multiplexing device (figure 4, structure 82) accessible by the processor over a fast bus (figure 4, combined structures of 80 and 22), the interrupt multiplexing device adapted to issue the interrupt to the processor in response to each pulsed interrupt signal generated by the external device (column 2, first paragraph); a low latency path (figure 4, dedicated bus between structures 18 and 82) coupling the external device to the interrupt multiplexing device and adapted to transport each pulsed interrupt signal generated by the external device to the interrupt multiplexing device; and a status bit (figure 6, structure 142) stored within the interrupt multiplexing device, the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device over the fast bus. Hence, claim is anticipated by the Greim.

Referring to claim 37: Greim discloses an external device (figure 4, structure 18 and the attached device from path 12) coupled to a high latency path (figure 4, structure 12), the external device generating a pulsed interrupt signal for each type of interrupt supported by the processor (figure 4, structure 20); an interrupt multiplexing device (figure 4, structure 82) accessible by the processor over a fast bus (figure 4, combined structures of 80 and 22), the interrupt multiplexing device adapted to issue the interrupt to the processor in response to each pulsed interrupt signal generated by the external device (column 2, first paragraph); a low latency path (figure 4, dedicated bus between structures 18 and 82) coupling the external device to the interrupt multiplexing device and adapted to transport each pulsed interrupt signal generated by the external device to the interrupt multiplexing device; and a status bit (figure 6, structure 142) stored within the interrupt multiplexing device, the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device over the fast bus. Hence, claim is anticipated by the Greim.

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

4. The factual inquiries set forth in *Graham* **v.** *John Deere Co.*, 383 U.S. 1, 148 USPQ 459 (1966), that are applied for establishing a background for determining obviousness under 35 U.S.C. 103(a) are summarized as follows:

1. Determining the scope and contents of the prior art.
2. Ascertaining the differences between the prior art and the claims at issue.
3. Resolving the level of ordinary skill in the pertinent art.
4. Considering objective evidence present in the application indicating obviousness or nonobviousness.

5. Claims 2-6 and 12-17 are rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of the Greim and Shek et al. (U.S. Patent No. 6,185,652).

Referring to claim 2: Greim's disclosure is stated above, but Greim does not explicitly disclose a current counter associated with the interrupt multiplexing device, the current counter incremented in response to each pulsed interrupt signal at the interrupt multiplexing device. Shek discloses a counter with the interrupt generator for counting the interrupt (abstract). Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt Shek's interrupt counter to Greim because Shek teaches one how to reduce the impact on the CPU power of maintaining an interrupt counter in the CPU (column 3, lines 51-52).

Referring to claim 3: Greim discloses that the processor is operable to service the stored interrupts and interrupt vectors in response to receiving the interrupt (column 2, lines 15-18). The means to service the interrupt is the claimed interrupt handler invoked by the processor.

Referring to claim 4: Shek discloses a last counter associated with the processor, the last counter incremented in response to each interrupt serviced by the CPU (column 4, lines 20-22, the count of tasks, column 4, lines 50-53, current count, task count, and current register).

Referring to claim 5: Shek discloses comparing a value of the last counter with a value of the current counter to determine whether there are more interrupts to service (column 4, lines 38-47).

Referring to claim 6: Greim's disclosure is stated above; furthermore, Greim discloses the DMA as the interrupt source (column 24, line 45).

Referring to claim 12: Shek discloses that it is known to initialize the counter (column 12, lines 29-30).

Referring to claim 13; Greim discloses reading the status bit; and if the status bit is clear, dismissing the handler (figure 8).

Referring to claim 14: Greim discloses clearing the status bit (figure 9).

Referring to claim 15: Shek discloses if the status bit is set, reading a value of a current counter; comparing the current counter value with a value of the last counter; and if the last counter value is greater than or equal to the current counter value, returning to the step of reading the status bit (column 4, lines 38-47). Shek discloses a last counter associated with the processor, the last counter incremented in response to each interrupt serviced by the CPU (column 4, lines 20-22, the count of tasks, column 4, lines 50-53, current count, task count, and current register).

Referring to claim 16: The admitted prior art discloses that determining whether the processor owns the control block is a part of well-known interrupt process.

Referring to claim 17: The admitted prior art discloses that it is a common-known interrupt process that when the processor owns the control block, it will process the control

block. Shek discloses that it is known to increment the counter when completing the interrupt (column 6, lines 55-57).

6. Claim 18 is rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of the Greim, Shek, and Ecclesine (U.S. Patent No. 5,983,275).

Referring to claim 18: Neither Greim nor Shek explicitly discloses determining whether a preset limit for processing control blocks has been reached; and if the preset limit is reached, dismissing the handler. Ecclesine discloses that it is known to determine whether a preset limit for processing control blocks has been reached for dismissing the handler (claim 38). Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt Ecclesine's teaching to Greim and Shek because Ecclesine teaches one to set a limit in distributing the interrupt handling processing power.

7. Claims 8-9 are rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of the Greim and Wu et al. (U.S. Patent No. 5,682,483) or Greim and the admitted prior art.

Referring to claims 8-9: Greim's disclosure is stated above, Greim does not explicitly disclose the printed circuit board trace and PCI. The admitted prior art discloses that the PCI is a well-known industrial practice (Application, page 3, first paragraph). Wu also discloses that the PCI is a popular industrial practice; in addition, Wu discloses that the PCI is implemented with the printed circuit board trace (column 5, lines 4-5). Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the

teachings of the admitted prior art or Wu because both the PCI and printed circuit board trace are well-known industrial practices.

8. Claim 10 is rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of the Greim and Shatas et al. (U.S. Patent No. 6,418,494).

Referring to claim 10: Greim's disclosure is stated above, Greim does not explicitly disclose the FPGA. Shatas discloses that it is known to employ a FPGA to distribute the communication signals among devices (figures 5-6). Shatas teaches that it is known to employ FPGA to distribute communication signals (figure 22). Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the teachings of the Shatas because Shatas teaches that it is known to use FPGA to distribute the communication signals.

9. Claims 1, 7, and 19-21 are rejected under 35 U.S.C. 103(a) as being unpatentable over the Swanstrom (U.S. Patent No. 5,754,884) in view of the Greim.

Referring to claim 1: Swanstrom discloses generating an interrupt signal at an external device (figure 1, structures 150 and 140), transporting the interrupt signal to an interrupt multiplexing device (figure 1, structure 160) over a first low latency path (figure 1, structures 142 and 152) that couples the external device to the interrupt multiplexing device, issuing the interrupt to the processor over a second low latency path (figure 1, structure 112). Swanstrom discloses a status register (figure 8, structure 812), but Swanstrom does not disclose that the register is within interrupt multiplexing device. Swanstrom also does not explicitly disclose

asserting a status bit in the interrupt multiplexing device corresponding to the interrupt in response to detecting the interrupt signal and reading the status bit over the second low latency path by an interrupt handler internal to the processor; clearing the status bit in response to the reading of the status bit to effectively acknowledge the interrupt. Swanstrom further does not disclose that the interrupt is a pulsed interrupt.

Greim discloses a status bit (figure 6, structure 142) stored within the interrupt multiplexing device (figure 6, structures 120, 124, and 142), the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device (abstract). Greim further discloses the edge sensitive interrupt (column 22, lines 35-36), which is the pulsed interrupt.

Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the Greim's teaching to Swanstrom because Greim teaches one how to reduce the impact on the processor and requestor's power by maintaining an status in the processor and requestor.

Referring to claim 7: claim 1's argument applies; furthermore, Swanstrom discloses a DMA controller (figure 1, structure 150, figure 7, structure 750).

Referring to claim 19: Swanstrom discloses generating an interrupt signal at an external device (figure 1, structures 150 and 140), transporting the interrupt signal to an interrupt multiplexing device (figure 1, structure 160) over a first low latency path (figure 1, structures 142 and 152) that couples the external device to the interrupt multiplexing device, issuing the interrupt to the processor over a second low latency path (figure 1, structure 112). Swanstrom

discloses a status register (figure 8, structure 812), but Swanstrom does not disclose that the register is within interrupt multiplexing device. Swanstrom also does not explicitly disclose asserting a status bit in the interrupt multiplexing device corresponding to the interrupt in response to detecting the interrupt signal and reading the status bit over the second low latency path by an interrupt handler internal to the processor; clearing the status bit in response to the reading of the status bit to effectively acknowledge the interrupt. Swanstrom further does not disclose that the interrupt is a pulsed interrupt.

Greim discloses a status bit (figure 6, structure 142) stored within the interrupt multiplexing device (figure 6, structures 120, 124, and 142), the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device (abstract). Greim further discloses the edge sensitive interrupt (column 22, lines 35-36), which is the pulsed interrupt.

Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the Greim's teaching to Swanstrom because Greim teaches one how to reduce the impact on the processor and requestor's power by maintaining an status in the processor and requestor.

Referring to claim 20: Swanstrom discloses generating an interrupt signal at an external device (figure 1, structures 150 and 140), transporting the interrupt signal to an interrupt multiplexing device (figure 1, structure 160) over a first low latency path (figure 1, structures 142 and 152) that couples the external device to the interrupt multiplexing device, issuing the interrupt to the processor over a second low latency path (figure 1, structure 112). Swanstrom

discloses a status register (figure 8, structure 812), but Swanstrom does not disclose that the register is within interrupt multiplexing device. Swanstrom also does not explicitly disclose asserting a status bit in the interrupt multiplexing device corresponding to the interrupt in response to detecting the interrupt signal and reading the status bit over the second low latency path by an interrupt handler internal to the processor; clearing the status bit in response to the reading of the status bit to effectively acknowledge the interrupt. Swanstrom further does not disclose that the interrupt is a pulsed interrupt.

Greim discloses a status bit (figure 6, structure 142) stored within the interrupt multiplexing device (figure 6, structures 120, 124, and 142), the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device (abstract). Greim further discloses the edge sensitive interrupt (column 22, lines 35-36), which is the pulsed interrupt.

Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the Greim's teaching to Swanstrom because Greim teaches one how to reduce the impact on the processor and requestor's power by maintaining an status in the processor and requestor.

Referring to claim 21: Swanstrom discloses generating an interrupt signal at an external device (figure 1, structures 150 and 140), transporting the interrupt signal to an interrupt multiplexing device (figure 1, structure 160) over a first low latency path (figure 1, structures 142 and 152) that couples the external device to the interrupt multiplexing device, issuing the interrupt to the processor over a second low latency path (figure 1, structure 112). Swanstrom

discloses a status register (figure 8, structure 812), but Swanstrom does not disclose that the register is within interrupt multiplexing device. Swanstrom also does not explicitly disclose asserting a status bit in the interrupt multiplexing device corresponding to the interrupt in response to detecting the interrupt signal and reading the status bit over the second low latency path by an interrupt handler internal to the processor; clearing the status bit in response to the reading of the status bit to effectively acknowledge the interrupt. Swanstrom further does not disclose that the interrupt is a pulsed interrupt.

Greim discloses a status bit (figure 6, structure 142) stored within the interrupt multiplexing device (figure 6, structures 120, 124, and 142), the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device (abstract). Greim further discloses the edge sensitive interrupt (column 22, lines 35-36), which is the pulsed interrupt.

Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the Greim's teaching to Swanstrom because Greim teaches one how to reduce the impact on the processor and requestor's power by maintaining an status in the processor and requestor.

10. Claims 2-6 and 12-17 are rejected under 35 U.S.C. 103(a) as being unpatentable over the Swanstrom in view of the Greim and Shek.

Referring to claim 2: Neither Swanstrom nor Greim explicitly disclose a current counter associated with the interrupt multiplexing device, the current counter incremented in response to

each pulsed interrupt signal at the interrupt multiplexing device. Shek discloses a counter with the interrupt generator for counting the interrupt (abstract). Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt Shek's interrupt counter to Greim because Shek teaches one how to reduce the impact on the CPU power of maintaining an interrupt counter in the CPU (column 3, lines 51-52).

Referring to claim 3: The interrupt handler invoked by the processor to service the issued interrupt is a well-known step in every interrupt process and is disclosed in the Application as the prior art.

Referring to claim 4: Shek discloses a last counter associated with the processor, the last counter incremented in response to each interrupt serviced by the CPU (column 4, lines 20-22, the count of tasks, column 4, lines 50-53, current count, task count, and current register).

Referring to claim 5: Shek discloses comparing a value of the last counter with a value of the current counter to determine whether there are more interrupts to service (column 4, lines 38-47).

Referring to claim 6: Greim's disclosure is stated above; furthermore, Greim discloses the DMA as the interrupt source (column 24, line 45).

Referring to claim 12: Shek discloses that it is known to initialize the counter (column 12, lines 29-30) and Shek discloses a last counter associated with the processor, the last counter incremented in response to each interrupt serviced by the CPU (column 4, lines 20-22, the count of tasks, column 4, lines 50-53, current count, task count, and current register).

Referring to claim 13; Greim discloses reading the status bit; and if the status bit is clear, dismissing the handler (figure 8).

Referring to claim 14: Greim discloses clearing the status bit (figure 9).

Referring to claim 15: Shek discloses if the status bit is set, reading a value of a current counter; comparing the current counter value with a value of the last counter; and if the last counter value is greater than or equal to the current counter value, returning to the step of reading the status bit (column 4, lines 38-47). Shek discloses a last counter associated with the processor, the last counter incremented in response to each interrupt serviced by the CPU (column 4, lines 20-22, the count of tasks, column 4, lines 50-53, current count, task count, and current register).

Referring to claim 16: The admitted prior art discloses that determining whether the processor owns the control block is a part of well-known interrupt process.

Referring to claim 17: The admitted prior art discloses that it is a common-known interrupt process that when the processor owns the control block, it will process the control block. Shek discloses that it is known to increment the counter when completing the interrupt (column 6, lines 55-57).

11. Claims 8-9 are rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of the Swanstrom, Greim and Wu or Swanstrom, Greim and the admitted prior art.

Referring to claims 8-9: Both Swanstrom and Greim's disclosures are stated above, neither explicitly discloses the printed circuit board trace and PCI. The admitted prior art discloses that the PCI is a well-known industrial practice (Application, page 3, first paragraph). Wu also discloses that the PCI is a popular industrial practice; in addition, Wu discloses that the PCI is implemented with the printed circuit board trace (column 5, lines 4-5). Hence, it would

have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the teachings of the admitted prior art or Wu to Swanstrom and Greim because both the PCI and printed circuit board trace are well-known industrial practices.

12. Claim 10 is rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of the Swanstrom, Greim and Shatas.

Referring to claim 10: Swanstrom and Greim's disclosures are stated above, neither explicitly discloses the FPGA. Shatas discloses that it is known to employ a FPGA to distribute the communication signals among devices (figures 5-6). Shatas teaches that it is known to employ FPGA to distribute communication signals (figure 22). Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the teachings of the Shatas to Swanstrom and Greim because Shatas teaches that it is known to use FPGA to distribute the communication signals.

13. Claim 18 is rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of the Swanstrom, Greim, Shek, and Ecclesine.

Referring to claim 18: Neither Greim nor Shek explicitly discloses determining whether a preset limit for processing control blocks has been reached; and if the preset limit is reached, dismissing the handler. Ecclesine discloses that it is known to determine whether a preset limit for processing control blocks has been reached for dismissing the handler (claim 38). Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant

made the invention to adapt Ecclesine's teaching to Swanstrom, Greim, and Shek because Ecclesine teaches one to set a limit in distributing the interrupt handling processing power.

14. Claims 22-24, 27-29, and 32-34 are rejected under 35 U.S.C. 103(a) as being unpatentable over the Swanstrom in view of the Greim and Okbay et al. (U.S. Patent No. 6,606,677).

Referring to claim 22: Swanstrom discloses generating an interrupt signal at an external device (figure 1, structures 150 and 140), transporting the interrupt signal to an interrupt multiplexing device (figure 1, structure 160) over a first low latency path (figure 1, structures 142 and 152) that couples the external device to the interrupt multiplexing device, issuing the interrupt to the processor over a second low latency path (figure 1, structure 112). Swanstrom discloses a status register (figure 8, structure 812), but Swanstrom does not disclose that the register is within interrupt multiplexing device. Swanstrom also does not explicitly disclose asserting a status bit in the interrupt multiplexing device corresponding to the interrupt in response to detecting the interrupt signal and reading the status bit over the second low latency path by an interrupt handler internal to the processor; clearing the status bit in response to the reading of the status bit to effectively acknowledge the interrupt. Swanstrom further does not disclose that the interrupt is a pulsed interrupt.

Greim discloses a status bit (figure 6, structure 142) stored within the interrupt multiplexing device (figure 6, structures 120, 124, and 142), the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing

device (abstract). Greim further discloses the edge sensitive interrupt (column 22, lines 35-36), which is the pulsed interrupt.

Neither Swanstrom nor Greim discloses that the interrupt handler is within the processor. Okbay discloses that it is known to integrate the interrupt handler within the processor to improve the interrupt handling (figure 3, structure 314).

Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the Greim and Okbay's teachings to Swanstrom because Greim teaches one how to reduce the impact on the processor and requestor's power by maintaining an status in the processor and requestor and Okbay teaches one to expedite the interrupt handling by integrating the interrupt handler within the processor.

Referring to claim 23: The admitted prior art discloses that determining whether the processor owns the control block and processing the control block are a part of well-known interrupt process.

Referring to claim 24: When the interrupt is completed and return back to the requester on the high latency path, it is the assigning ownership of the control block over a high latency path.

Referring to claim 27: Swanstrom discloses generating an interrupt signal at an external device (figure 1, structures 150 and 140), transporting the interrupt signal to an interrupt multiplexing device (figure 1, structure 160) over a first low latency path (figure 1, structures 142 and 152) that couples the external device to the interrupt multiplexing device, issuing the interrupt to the processor over a second low latency path (figure 1, structure 112). Swanstrom discloses a status register (figure 8, structure 812), but Swanstrom does not disclose that the

register is within interrupt multiplexing device. Swanstrom also does not explicitly disclose asserting a status bit in the interrupt multiplexing device corresponding to the interrupt in response to detecting the interrupt signal and reading the status bit over the second low latency path by an interrupt handler internal to the processor; clearing the status bit in response to the reading of the status bit to effectively acknowledge the interrupt. Swanstrom further does not disclose that the interrupt is a pulsed interrupt.

Greim discloses a status bit (figure 6, structure 142) stored within the interrupt multiplexing device (figure 6, structures 120, 124, and 142), the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device (abstract). Greim further discloses the edge sensitive interrupt (column 22, lines 35-36), which is the pulsed interrupt.

Neither Swanstrom nor Greim discloses that the interrupt handler is within the processor. Okbay discloses that it is known to integrate the interrupt handler within the processor to improve the interrupt handling (figure 3, structure 314).

Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the Greim and Okbay's teachings to Swanstrom because Greim teaches one how to reduce the impact on the processor and requestor's power by maintaining an status in the processor and requestor and Okbay teaches one to expedite the interrupt handling by integrating the interrupt handler within the processor.

Referring to claim 28: The admitted prior art discloses that determining whether the processor owns the control block and processing the control block are a part of well-known interrupt process.

Referring to claim 29: When the interrupt is completed and return back to the requester on the high latency path, it is the assigning ownership of the control block over a high latency path.

Referring to claim 32: Swanstrom discloses generating an interrupt signal at an external device (figure 1, structures 150 and 140), transporting the interrupt signal to an interrupt multiplexing device (figure 1, structure 160) over a first low latency path (figure 1, structures 142 and 152) that couples the external device to the interrupt multiplexing device, issuing the interrupt to the processor over a second low latency path (figure 1, structure 112). Swanstrom discloses a status register (figure 8, structure 812), but Swanstrom does not disclose that the register is within interrupt multiplexing device. Swanstrom also does not explicitly disclose asserting a status bit in the interrupt multiplexing device corresponding to the interrupt in response to detecting the interrupt signal and reading the status bit over the second low latency path by an interrupt handler internal to the processor; clearing the status bit in response to the reading of the status bit to effectively acknowledge the interrupt. Swanstrom further does not disclose that the interrupt is a pulsed interrupt.

Greim discloses a status bit (figure 6, structure 142) stored within the interrupt multiplexing device (figure 6, structures 120, 124, and 142), the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing

device (abstract). Greim further discloses the edge sensitive interrupt (column 22, lines 35-36), which is the pulsed interrupt.

Neither Swanstrom nor Greim discloses that the interrupt handler is within the processor. Okbay discloses that it is known to integrate the interrupt handler within the processor to improve the interrupt handling (figure 3, structure 314).

Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the Greim and Okbay's teachings to Swanstrom because Greim teaches one how to reduce the impact on the processor and requestor's power by maintaining an status in the processor and requestor and Okbay teaches one to expedite the interrupt handling by integrating the interrupt handler within the processor.

Referring to claim 33: Swanstrom discloses generating an interrupt signal at an external device (figure 1, structures 150 and 140), transporting the interrupt signal to an interrupt multiplexing device (figure 1, structure 160) over a first low latency path (figure 1, structures 142 and 152) that couples the external device to the interrupt multiplexing device, issuing the interrupt to the processor over a second low latency path (figure 1, structure 112). Swanstrom discloses a status register (figure 8, structure 812), but Swanstrom does not disclose that the register is within interrupt multiplexing device. Swanstrom also does not explicitly disclose asserting a status bit in the interrupt multiplexing device corresponding to the interrupt in response to detecting the interrupt signal and reading the status bit over the second low latency path by an interrupt handler internal to the processor; clearing the status bit in response to the reading of the status bit to effectively acknowledge the interrupt. Swanstrom further does not disclose that the interrupt is a pulsed interrupt.

Greim discloses a status bit (figure 6, structure 142) stored within the interrupt multiplexing device (figure 6, structures 120, 124, and 142), the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device (abstract). Greim further discloses the edge sensitive interrupt (column 22, lines 35-36), which is the pulsed interrupt.

Neither Swanstrom nor Greim discloses that the interrupt handler is within the processor. Okbay discloses that it is known to integrate the interrupt handler within the processor to improve the interrupt handling (figure 3, structure 314).

Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the Greim and Okbay's teachings to Swanstrom because Greim teaches one how to reduce the impact on the processor and requestor's power by maintaining an status in the processor and requestor and Okbay teaches one to expedite the interrupt handling by integrating the interrupt handler within the processor.

Referring to claim 34: Swanstrom discloses generating an interrupt signal at an external device (figure 1, structures 150 and 140), transporting the interrupt signal to an interrupt multiplexing device (figure 1, structure 160) over a first low latency path (figure 1, structures 142 and 152) that couples the external device to the interrupt multiplexing device, issuing the interrupt to the processor over a second low latency path (figure 1, structure 112). Swanstrom discloses a status register (figure 8, structure 812), but Swanstrom does not disclose that the register is within interrupt multiplexing device. Swanstrom also does not explicitly disclose asserting a status bit in the interrupt multiplexing device corresponding to the interrupt in

response to detecting the interrupt signal and reading the status bit over the second low latency path by an interrupt handler internal to the processor; clearing the status bit in response to the reading of the status bit to effectively acknowledge the interrupt. Swanstrom further does not disclose that the interrupt is a pulsed interrupt.

Greim discloses a status bit (figure 6, structure 142) stored within the interrupt multiplexing device (figure 6, structures 120, 124, and 142), the status bit adapted for assertion whenever the pulsed interrupt signal is detected at the interrupt multiplexing device, wherein the processor efficiently acknowledges the issued interrupt by accessing the interrupt multiplexing device (abstract). Greim further discloses the edge sensitive interrupt (column 22, lines 35-36), which is the pulsed interrupt.

Neither Swanstrom nor Greim discloses that the interrupt handler is within the processor. Okbay discloses that it is known to integrate the interrupt handler within the processor to improve the interrupt handling (figure 3, structure 314).

Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt the Greim and Okbay's teachings to Swanstrom because Greim teaches one how to reduce the impact on the processor and requestor's power by maintaining an status in the processor and requestor and Okbay teaches one to expedite the interrupt handling by integrating the interrupt handler within the processor.

15. Claims 25 and 30 are rejected under 35 U.S.C. 103(a) as being unpatentable over the Swanstrom in view of the Greim, Okbay, and Shek.

Referring to claims 25 and 30: The disclosures of Swanstrom, Greim, and Okbay are stated above. None of them explicitly discloses a current counter associated with the interrupt multiplexing device, the current counter incremented in response to each pulsed interrupt signal at the interrupt multiplexing device. Shek discloses a counter with the interrupt generator for counting the interrupt (abstract). Shek further discloses a last counter associated with the processor, the last counter incremented in response to each interrupt serviced by the CPU (column 4, lines 20-22, the count of tasks, column 4, lines 50-53, current count, task count, and current register). Shek also discloses comparing a value of the last counter with a value of the current counter to determine whether there are more interrupts to service (column 4, lines 38-47).

Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt Shek's interrupt counter to Swanstrom, Greim, and Okbay because Shek teaches one how to reduce the impact on the CPU power of maintaining an interrupt counter in the CPU (column 3, lines 51-52).

16. Claim 26 and 31 are rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of the Swanstrom, Greim, Okbay, Shek, and Ecclesine.

Referring to claim 26 and 31: None of Swanstrom, Greim, Okbay, and Shek explicitly discloses determining whether a preset limit for processing control blocks has been reached; and if the preset limit is reached, dismissing the handler. Ecclesine discloses that it is known to determine whether a preset limit for processing control blocks has been reached for dismissing the handler (claim 38). Hence, it would have been obvious to one having ordinary skill in the computer art at the time Applicant made the invention to adapt Ecclesine's teaching to

Swanstrom, Greim, Okbay, and Shek because Ecclesine teaches one to set a limit in distributing the interrupt handling processing power.

Response to Arguments

17. In response to Applicant's argument that the prior arts on record do not disclose the claimed limitations (Remark, page 16, 2nd paragraph, page 17, 2nd paragraph, page 19, 3rd paragraph): As the aforementioned rejections indicate, the prior arts on record do disclose the claimed limitations. Greim discloses "**an interrupt multiplexing device** (figure 4, structure 82) **accessible by the processor** (figure 4, structure 20) **over a fast bus** (figure 4, combined structures of 80 and 22)", and **"a low latency path** (figure 4, the dedicated bus between structures 18 and 82) **coupling the external device** (figure 4, structure 18 and the attached device from path 12) **to the interrupt multiplexing device".** Greim discloses that the interrupt controls includes a receiving device and a processor interrupt generating device (column 2, lines 7-8, 12-14, and 22-27, and claim 1), and the processor will send information to the processor interrupt generating device for acknowledge (column 2, lines 27-30). Since Greim's processor (figure 4, structure 20) directly connects to the interrupt controller (figure 4, structure 82) without accessing the low latency bus (figure 4, structure 12), thus, Greim discloses **"the processor acknowledges the issued interrupt without accessing the high latency path by accessing the interrupt multiplexing device over the fast bus".**

Swanstrom discloses "**an interrupt multiplexing device** (figure 1, structure 160) **accessible by the processor** (figure 1, structure 110) **over a fast bus** (figure 1, structure 112)",

and **"a low latency path (**figure 1, structure 152**) coupling the external device (**figure 1, structure 150**) to the interrupt multiplexing device".**

18. In response to Applicant's argument that the prior art Greim issues interrupt acknowledge signals to the external device, rather to the interrupt controller as claimed in the Application (Remark, page 16, 3rd paragraph, lines 3-4, page 17, first paragraph): Greim does disclose issuing acknowledge to the interrupt controller. Greim discloses that the interrupt controls includes a receiving device and a processor interrupt generating device (column 2, lines 7-8, 12-14, and 22-27, and claim 1), and the processor will send information to the processor interrupt generating device for acknowledge (column 2, lines 27-30).

Conclusion

19. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Justin I. King whose telephone number is 703-305-4571. The examiner can normally be reached on Monday through Friday, 9:00 am to 5:00 pm. If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Mark Rinehart can be reached on 703-308-3110. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR

system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Lastly, paper copies of cited U.S. patents and U.S. patent application publications will cease to be mailed to applicants with Office actions as of June 2004. Paper copies of foreign patents and non-patent literature will continue to be included with office actions. These cited U.S. patents and patent application publications are available for download via the Office's PAIR. As an alternate source, all U.S. patents and patent application publications are available on the USPTO web site (www.uspto.gov), from the Office of Public Records and from commercial sources. Applicants are referred to the Electronic Business Center (EBC) at http://www.uspto.gov/ebc/index.html or 1-866-217-9197 for information on this policy. Requests to restart a period for response due to a missing U.S. patent or patent application publications will not be granted.

MARK H. RINEHART
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100

Justin King
September 2, 2004